

When a Fleet Needs a New Skill Before Monday

A fleet reliability lead has until Monday's opening inbound window to teach forty-one mobile handling units one new behavior, and the thing she cannot do is find out in advance what that behavior will do to everything the units already know. Her deadline is a customer's qualification week, and a qualification week is observed once. This article describes, as a technical matter, the governance-credentialed model-adaptation primitive disclosed in U.S. Provisional Application No. 64/049,409: adaptation artifacts carrying capability scope, dependency, certification, and lineage fields; sandbox evaluation performed before activation on the operational substrate; composition that produces a fresh certification requirement; and deactivation and rollback pathways recorded in the consuming agent's lineage field.

Friday, 4:10 p.m., Forty-One Units and One New Film

The fleet reliability lead at a regional cold-storage logistics operator has forty-one mobile handling units working three docks, and a new account whose opening inbound trailers arrive at 5:00 a.m. Monday. She learned on Thursday that this customer's pallets ship under a slick, low-friction stretch film that her units read as a different

surface class than anything on the docks today. On the two test pallets she ran that afternoon, three units out of five approached correctly and set down early, and two adjusted grip in a way that left the top cases proud of the stack.

She has a candidate handling skill. A supplier sent it that morning, and it plausibly covers exactly this film. What she does not have, at 4:10 on Friday, is any way to answer the three questions she would need answered before putting it on the fleet. She cannot tell what the candidate does to the pick-and-place behavior her units already run, tuned over eleven months against the shrink-wrapped and banded loads that make up the rest of her volume. She cannot tell what it does when a unit is outside the narrow situation it was built for, on a mixed trailer at 5:40 a.m. where two pallets in nine carry the new film. And she cannot tell whether, if a unit misbehaves at 6:15 Monday with the dock full, she can get the candidate back off that unit, and off the other forty, without walking to each one.

So she does what her setup leaves her: she loads nothing, she staffs the Monday inbound with four manual handlers she has to call at home on a Friday evening, and she tells the account manager that week one will be slower than the pitch said.

What Monday Costs Her, and Why She Cannot Run It Again

The manual Monday is not the loss. It is what she does instead, and it costs her a weekend of calls and a week of unbudgeted overtime.

The loss is that the customer's opening week is the qualification week. Their inbound quality team is on site Monday through Thursday, sampling receipts, timing dock turns, and writing the observations that become the service baseline in the contract's operating annex. Whatever her docks do while those people are standing on them becomes the number she is measured against for the term, and the throughput she quoted assumed automated handling of that film from trailer one. She will not be given

a second qualification week because her fleet was not ready for the one she got, and the annex will not be reopened on the argument that it would have gone better in week three.

There is a second, quieter loss in the same decision. Every week she runs the new film manually is a week her units generate no observations against it, so the operating history she would want before trusting any handling change on that surface does not accumulate. Her caution on Friday is what keeps her cautious in April.

Had she gone the other way, the loss would change shape. Loading an unevaluated behavior on all forty-one units before a qualification week puts her eleven months of tuned handling at risk in front of the audience whose opinion is contractually fixed. A dropped case at 6:15 Monday is product, cold chain, and a written observation, in that order.

Why Her Friday Has No Good Move

Her problem is not that she doubts the supplier. It is that in her setup the artifact she received and the knowledge she would need about it live in different places, and in her setup the second one does not travel to her.

The file that arrived Thursday tells her what it is called and roughly what it is for. Nothing in the package she received states the boundary of the situations it was built to cover, so she has no machine-checkable way to say that a unit running it on a banded pallet is outside its intended range. Nothing attached to it says which of her currently loaded behaviors it assumes are also running, so if the supplier tuned it on top of a baseline her units do not have, she would find that out on the dock. Nothing attached to it records what it was trained on or who stands behind it in a form her fleet could evaluate before activation, so the trust decision is one she makes at a desk, once, for all forty-one units.

That is the compounding part. Because her fleet has no way to evaluate the candidate against its own current state, her granularity is the whole fleet and her test environment is Monday. Were her units able to run the candidate against a copy of their own configuration before it touched the operational one, Friday would be a different afternoon: she would be reading a conflict report instead of guessing at one. As her deployment is configured today, she has the choice between shipping an unknown into a qualification week and not shipping at all, and both are decisions about the account rather than about the software.

Reversal is the piece that makes her pick the cautious side every time. If getting a behavior back off her units means touching them one at a time in a hot dock, then for her purposes activation is close enough to permanent that she treats it that way, and she declines updates she would otherwise want. Her fleet gets more conservative every quarter for a reason that has nothing to do with whether the updates were any good.

What the Filed Mechanism Does

U.S. Provisional Application No. 64/049,409 discloses, in Chapter 8, a governance-credentialed model-adaptation primitive by which governed agents extend a base cognitive substrate through composable adaptation artifacts. In an embodiment, each adaptation artifact is itself a governance-credentialed governed observation carrying an authority credential and lineage-bound provenance, distributed through the governed mesh, composed at inference time, and subject to the composite admissibility evaluator at each composition, activation, and evaluation.

The artifact primitive of Section 8.2, laid out as a byte sequence in FIG. 8A, carries in an embodiment fields of the kind her Friday file lacks. Alongside the artifact identifier 800a, the adaptation-technique identifier 800b, and the artifact-content field 800c, an embodiment carries a capability scope specification 800d identifying the bounded range of cognitive contexts within which the artifact is governance-policy-defined to apply; a compatibility specification 800e; a dependency specification 800g identifying

prerequisite artifacts; a certification record 800h; a provenance-lineage record 800i identifying training inputs, training methodology, and contributing authorities; an authority credential 800j; a version identifier and version-lineage chain 800l; and a cryptographic integrity attestation 800m over the preceding fields.

Section 8.5 addresses the test environment she does not have. In an embodiment, a sandbox evaluation mechanism evaluates a candidate artifact in a governance-credentialed sandboxed execution environment prior to activation on the consuming agent's operational cognitive substrate. It comprises a sandboxed-substrate isolator that instantiates a sandboxed copy of that agent's substrate separate from the operational one; an evaluation-context generator drawing representative contexts from the agent's current task context, from governance-policy-defined standard evaluation suites, and from the agent's experiential observation store; a composition executor applying the candidate; a behavior evaluator producing a compatibility assessment identifying conflicts with currently loaded artifacts on the operational substrate; a performance preview generator; a governance summary generator covering licensing constraints, capability-scope boundaries, dependency satisfaction, and authority compatibility; a risk-projection evaluator; and an activation-gate controller producing an admit-or-reject outcome. The resulting certification record is lineage-recorded and, in an embodiment, is admissible as input to admissibility evaluations at other consuming agents considering the same artifact. Where an initial fast evaluation returns an outcome with elevated uncertainty, an embodiment escalates to a deeper subsequent evaluation with expanded contexts, longer duration, and stricter criteria.

Section 8.4 addresses the stack rather than the single artifact. In an embodiment, artifacts compose through patterns including simultaneous composition with weighted combination at inference time, sequential swapping with governance-policy-defined handoff semantics, offline merging, hierarchical composition with declared dependencies, contextual composition, and ensemble composition. In an embodiment, each governed composition event produces a composite artifact that is itself an adaptation artifact, inheriting a composite capability scope equal to the intersection of

source scopes, or their union where governance-policy-defined, a merged provenance-lineage record preserving each source's identity and proportional contribution, and a fresh certification record requiring sandbox evaluation of the composite prior to activation.

Section 8.10 addresses the mixed trailer. In an embodiment, the composite admissibility evaluator performs a governance gate and a skill-domain-routing function at once: on determining that a candidate generation step falls outside the capability scope of all active artifacts, it triggers a routing response, which per the disclosure may include activating a different already-loaded artifact whose scope covers the current context, reducing the confidence value to reflect generation outside governed competence boundaries with graduated-actuation mode de-escalation per Section 6.4, deferring pending a governed discovery query, continuing in degraded mode under stricter criteria, engaging consultative mode requesting human-operator input, or rolling generation back to a prior checkpoint. Each routing decision is lineage-recorded.

Reversal is addressed in Sections 8.9 and 8.11. In an embodiment, a cascade-deactivation mechanism deactivates dependent artifacts on deactivation of a strict prerequisite; a deprecation observation signals consuming agents to transition to a successor within a governance-policy-defined transition window, with soft, hard, security-driven, and regulatory-driven deprecation patterns available; and a rollback mechanism returns a consuming agent to a prior-version artifact on detection of post-activation governance-policy violations attributable to the current version, recording the rollback event, the triggering violation, and the restored artifact in the agent's lineage field.

Still on Her Desk Afterward

The disclosed architecture would not have written her handling skill. Whether an artifact that actually grips that film well exists, and whether the supplier's version is any good, are questions about the artifact's content, and no field in the layout answers

them for her.

Nor would the outcome of a sandbox run be better than the contexts it drew on. In her deployment the evaluation contexts would come from her units' own task context, her own observation store, and whatever standard suites her governance policy names, so a failure mode her docks have not produced is one her sandbox would have no sample of. The disclosure conditions the pre-certification shortcut too: her fleet's ability to admit somebody else's certification instead of running its own turns on pre-certification-trust parameters that her deployment would have to set, and setting them is her judgment call, not the architecture's.

Deactivation, as disclosed, reaches artifacts. It would not undo a case her unit already dropped at 6:15, or unwrite the observation the customer's quality team made while standing there. And none of the credentialing described above would tell her whether the authority that signed the artifact deserved the standing it claims; the mechanisms make the credential, the scope, and the lineage available to her fleet's evaluation, and her governance policy still has to say what to do with them.

Disclosure Scope

This article is a technical description of subject matter disclosed in U.S. Provisional Application No. 64/049,409, including the governance-credentialed model-adaptation primitive, the adaptation artifact primitive and its structural fields, the governed composition mechanism, the sandbox and pre-activation evaluation mechanism, the dependency and prerequisite chain mechanism with cascade deactivation, the admissibility-gate-as-router mechanism, and the adapter lifecycle, versioning, and deprecation mechanism described therein. The scenario, party, and deployment described above are illustrative and are not drawn from that filing.

Nothing in this article characterizes the scope of any claim, present or future, in that filing or in any application claiming its benefit, and nothing here is an admission regarding the state of the art. Statements about outcomes follow the filing's own conditioning: where the disclosure conditions a result on a governance policy, a declared bound, a credential, or an evaluation outcome, that condition is stated here as well.

Spatial Adaptation Artifacts (</spatial-adaptation>) [All 49 steps → \(/inventive-steps\)](#)

Runtime signed-skill loading. Sandbox-certified. Admissibility gate as skill router.

Provisional application

PRIMARY TECHNICAL DISCLOSURE

- [Spatial Adaptation Artifacts: Runtime Skill Loading With Admissibility Gating \(/articles/spatial-adaptation-artifacts-runtime-skill-loading-with-admissibility-gating\)](/articles/spatial-adaptation-artifacts-runtime-skill-loading-with-admissibility-gating)

SECONDARY TECHNICAL

- [Runtime-Signed Adaptation Artifacts \(/articles/spatial-adaptation/runtime-signed-artifacts\)](/articles/spatial-adaptation/runtime-signed-artifacts)
- [Sandbox Pre-Activation Certification \(/articles/spatial-adaptation/sandbox-pre-activation-certification\)](/articles/spatial-adaptation/sandbox-pre-activation-certification)
- [Admissibility as Skill Router \(/articles/spatial-adaptation/admissibility-as-skill-router\)](/articles/spatial-adaptation/admissibility-as-skill-router)
- [Always-Active Personal Layer \(/articles/spatial-adaptation/always-active-personal-layer\)](/articles/spatial-adaptation/always-active-personal-layer)
- [Cascade Deactivation Dependencies \(/articles/spatial-adaptation/cascade-deactivation-dependencies\)](/articles/spatial-adaptation/cascade-deactivation-dependencies)
- [Cross-Model Adaptation Portability \(/articles/spatial-adaptation/cross-model-portability\)](/articles/spatial-adaptation/cross-model-portability)
- [Federated Skill Training \(/articles/spatial-adaptation/federated-skill-training\)](/articles/spatial-adaptation/federated-skill-training)
- [Composite Licensing Intersection \(/articles/spatial-adaptation/composite-licensing-intersection\)](/articles/spatial-adaptation/composite-licensing-intersection)
- [Decentralized Mesh Adaptation Distribution \(/articles/spatial-adaptation/decentralized-mesh-distribution\)](/articles/spatial-adaptation/decentralized-mesh-distribution)

APPLICATIONS · GENERAL

- [Governed In-Field AI Model Adaptation for Defense Operations \(/articles/spatial-adaptation/defense-field-adaptation\)](/articles/spatial-adaptation/defense-field-adaptation)
- [Safe Runtime Updates for AI-Driven Industrial Robots: Governed Adaptation Under ISO 10218 and the EU Machinery Regulation \(/articles/spatial-adaptation/industrial-robotics-adaptive-update\)](/articles/spatial-adaptation/industrial-robotics-adaptive-update)
- [Governing Adaptive Medical Device and SaMD Updates Under FDA PCCP and EU MDR \(/articles/spatial-adaptation/medical-device-adaptive-update\)](/articles/spatial-adaptation/medical-device-adaptive-update)
- [Safe Rapid Security Updates for Safety-Critical Systems \(/articles/spatial-adaptation/cybersecurity-rapid-update\)](/articles/spatial-adaptation/cybersecurity-rapid-update)
- [Regulatory-Aware LLM Adaptation: Verifiable Governance for EU AI Act and FDA Compliance \(/articles/spatial-adaptation/regulatory-aware-llm-adaptation\)](/articles/spatial-adaptation/regulatory-aware-llm-adaptation)
- [IEC 62304 Compliance for Continuously Adapting Medical Device Software \(/articles/spatial-adaptation/iec-62304-medical-software\)](/articles/spatial-adaptation/iec-62304-medical-software)
- [Enforcing NIST AI RMF Compliance at Runtime for AI Model Fleets \(/articles/spatial-adaptation/nist-ai-rmf\)](/articles/spatial-adaptation/nist-ai-rmf)
- [UNECE R156 SUMS Compliance for In-Vehicle Software Updates \(/articles/spatial-adaptation/unec-r156-software-update\)](/articles/spatial-adaptation/unec-r156-software-update)
- [When a Fleet Needs a New Skill Before Monday \(/articles/spatial-adaptation/spatial-adaptation-stakes\)](/articles/spatial-adaptation/spatial-adaptation-stakes)

APPLICATIONS · SPECIFIC

- [Governed Adaptation Beyond Anthropic Skills: Admissibility-Gated, Reversible Capability Loading \(/articles/spatial-adaptation/anthropic-skills\)](/articles/spatial-adaptation/anthropic-skills)
- [OpenAI Fine-Tuning vs Governed Model Adaptation \(/articles/spatial-adaptation/openai-fine-tuning\)](/articles/spatial-adaptation/openai-fine-tuning)
- [Tesla FSD Updates vs Governed Adaptation Artifacts \(/articles/spatial-adaptation/tesla-fsd-updates\)](/articles/spatial-adaptation/tesla-fsd-updates)
- [AWS Bedrock vs Governed Adaptation: The Certifiable-Artifact Layer \(/articles/spatial-adaptation/aws-bedrock\)](/articles/spatial-adaptation/aws-bedrock)
- [Databricks Mosaic AI vs Governed Adaptation Artifacts \(/articles/spatial-adaptation/databricks-ai\)](/articles/spatial-adaptation/databricks-ai)
- [Google Vertex AI vs. Governed Adaptation: Who Owns Model-Adaptation Governance? \(/articles/spatial-adaptation/google-vertex-ai\)](/articles/spatial-adaptation/google-vertex-ai)
- [Hugging Face Hub Alternative for Governed Model Adaptation \(/articles/spatial-adaptation/huggingface-hub\)](/articles/spatial-adaptation/huggingface-hub)
- [Governed Agent Adaptation vs Anthropic Claude MCP \(/articles/spatial-adaptation/anthropic-claude-mcp\)](/articles/spatial-adaptation/anthropic-claude-mcp)

HOW-TO GUIDES

- [How to Meet Regulatory Software-Update Rules for Autonomous Vehicles \(/articles/guides/meet-vehicle-software-update-regulations\)](/articles/guides/meet-vehicle-software-update-regulations)
- [How to Safely Push Field Updates to Autonomous Systems \(/articles/guides/safely-push-field-updates-to-autonomous-systems\)](/articles/guides/safely-push-field-updates-to-autonomous-systems)
- [How to Safely Roll Back a Bad Update to an Autonomous Fleet \(/articles/guides/safely-roll-back-a-n-autonomous-fleet-update\)](/articles/guides/safely-roll-back-a-n-autonomous-fleet-update)

[Spatial Adaptation Artifacts overview → \(/spatial-adaptation\)](/spatial-adaptation)